Robotics Research Technical Report

Generatorium omnis laboris ex machina

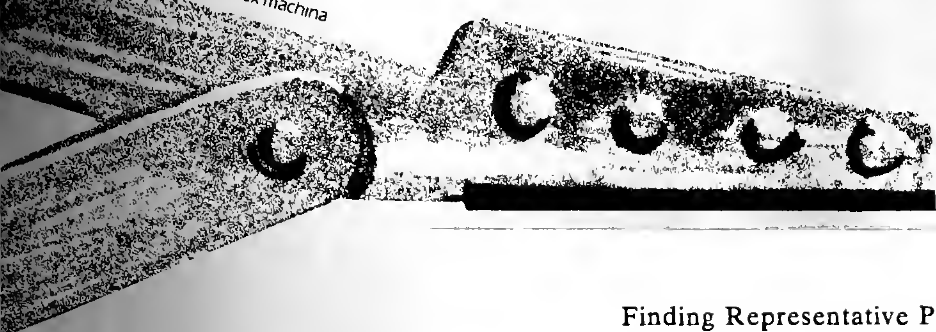

Finding Representative Points of Closest Approach for Noisy Curves

by

C. Marc Bastuscheck

Technical Report No. 250
Robotics Report No. 83
October, 1986

New York University
Co nstitute of Mathematical Sciences
Computer Science Division
51 Mercer Street New York, N.Y. 10012

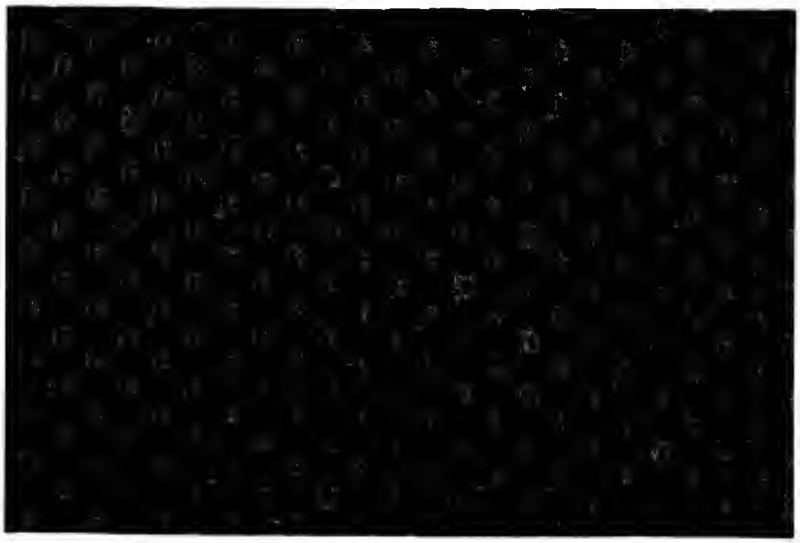

Finding Representative Points of
Closest Approach for Noisy Curves

by

C. Marc Bastuscheck

Technical Report No. 250
Robotics Report No. 83
October, 1986

New York University
Dept. of Computer Science
Courant Institute of Mathematical Sciences
251 Mercer Street
New York, New York 10012

Work on this paper has been supported by Office of Naval Research Grant N00014-82-K-0381, National Science Foundation CER Grant DCR-83-20085, and by grants from the Digital Equipment Corporation and the IBM Corporation.

Finding Representative Points of Closest Approach for Noisy Curves

C. Marc Bastuscheck

Robotics Activity
Courant Institute
New York University

ABSTRACT

Matching two sets of curves in three-space can be accomplished efficiently when small sets of points are used to represent each set. Two methods of choosing representative 'points of closest approach' of a pair of smooth curves in three-space which have been corrupted by noise are investigated. The methods are 1) aggregation of a centroid-like point, and 2) determination of points of closest approach along polynomial fits to the noisy curves in the region of closest approach. Random noise with a pseudo-normal distribution was added to computer generated curves to provide an environment of controlled data for testing the different point selection procedures. Performance of the procedures was evaluated by measuring the separation of two sets of curves after they had been matched using representative points. It was observed that the 'centroid' method allows better selection of points than quadratic or cubic fits when substantial lengths of the curves can be used, but that a cubic fit of coordinates vs arc length gave better results when relatively short lengths of curve were used. The quadratic fits behaved very badly. The results have application to data reduction for efficient recognition of three dimensional objects, and the routine for measuring separation between sets of curves has many interesting applications.

1. Introduction

An approach which may prove useful in object recognition by computer is the extraction of significant curves from the mass of low level vision data made available, for example, by an imaging range sensor. Object recognition would then reduce to a process of matching sets of curves in three-space. If a small set of points could be found which would effectively represent a set of curves object recognition could be performed even more efficiently Such a set of points might be defined in many ways, only one of which is considered here: for some pairs of curves points of closest approach are well defined and can be exploited if such points can be easily defined for noisy curves.

The following problem is considered: given a pair of smooth curves in three-space which have well defined points of closest approach, how can these points be best determined when the curves are corrupted by noise. By way of illustration, the dashed line segments in Fig. 1 have as endpoints some of the points of closest approach for a set of curves. In this work all curves have been chosen so each pair of curves has at most one pair of points of closest approach. Since object recognition lies at the heart of this work, success of procedures to generate representative

Work on this paper has been supported by Office of Naval Research Grant N00014-82-K-0381, National Science Foundation CER Grant DCR-83-20085, and by grants from the Digital Equipment Corporation and the IBM Corporation

points is measured by the quality of matches made using the representative points.

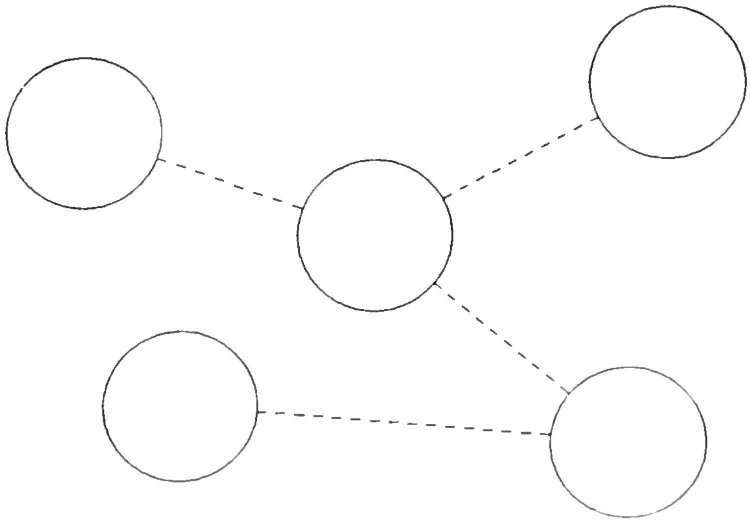

Fig. 1 Dashed lines connect some "points of closest approach"

Points representative of curves such as the circles in Fig. 1 could be defined in many ways. For example, centroids of the curves might serve; however, if curves are less than fully observed the centroid of a partly observed object might not correspond properly with the centroid of a fully observed model. Points of closest approach overcome this difficulty by providing a set of points which usefully describes a set of curves and consists of points which are independent of translation, rotation, and partial reduction of the data set. If such representative points are well chosen they will also be relatively unaffected by noise. The experiments reported here explore two methods by which representative points of closest approach can be generated with particular attention to the effects of random noise.

The paper is organized as follows. Section II provides background information on work in object recognition and explanation helpful in understanding these experiments. In Section III two methods for choosing representative points of closest approach are described, and in addition procedures used to generate synthetic "depth" measurements of objects, filters to degrade perfect data through the addition of random noise and other distortions, a routine to match two sets of points and a procedure for measuring the spatial separation of two sets of lines in three-space are given. In Section IV the behavior of the two different point generation methods is examined closely for different conditions of noise and amount of length of arc use. In section V the results are summarized and considered in the broader context of robotics.

2. Background and Motivation

Computer recognition of an object can be considered to be comparison of an observed object with a set of model objects to discover the identity, location, and orientation of the observed object. A key element of any successful procedure for object recognition must be reduction of the enormous amount of raw vision data to an object description which can be easily compared with a number of models to find a matching object. Reduction of depth measurements (or depth-related images, such as those used in photometric stereo) has often consisted of generating regions (usually planar) whose character and connectivity can be used for object identification [SS71][OS81][FH83][B84][H85]. Bolles and coworkers used clusters of features such as holes and corners to determine the identification and location of 2D industrial parts from 2D images [BC82], and extended this work by extracting edges and planes as features from range images to determine the identity and location of 3D parts [HB84]. The recent review by Besl and Jain [BJ85] covers many aspects of the problem of object recognition.

One method of object recognition proceeds from the assumption that suitable lines in three space can be used effectively to describe 3d objects. Suitable lines might be edges (of various types, such as sharp discontinuities, occluding contours, changes in reflectivity on a smooth surface) or geometrically defined lines (such as the line where a cylinder becomes tangent to a plane). There should be no requirement that lines be easily parametrized, although special cases of straight lines and circles are likely to be of considerable interest in the identification of machined objects.

A method for describing a general curve in three-space has been given [BSSS86] which uses a list of points, and such descriptions were used to identify an observed curve as being the same as or different from a previously observed curve. Extension of that matching procedure to objects described by more than one curve is not obvious, since curves from observations of an object in different orientations may be recorded in different orders, and matching all sub-curves with all others would increase computation combinatorially.

Computation can be significantly reduced if a small set of representative points is used to describe an entire object. Representative points might be vertices or centers of holes, and many other types of representative points can be imagined. In the present work objects were described by sets of curves which (in the absence of noise) would have unique and useful points of closest approach. While not equivalent to a full description of the curves, a sub-set of these points of closest approach could be used as geometric features, since these points are independent of the location and orientation of the object and can be used to match points generated in a similar manner from a model set of curves. When selected points of object and model are given in the same order the computation required for matching is quite small.

Matching is performed using the same least squares matching algorithm of [BSSS86], but only a very small number of points is used. This routine returns the translation and rotation which must be applied to an observed object to bring it into best juxtaposition with a model, and the least squares distance between the sets of points matched. However, this distance cannot be used as a measure of the quality of the match except when significant mismatch is indicated. If representative points are badly chosen (e.g. nearly colinear) or if they are affected by noise the matching routine may return a distance indicating a good match while the overlay of the observed object on the model looks terrible. For this reason a method was developed to measure the separation between two sets of curves which are supposed to be close in real space.

When two curves are described with infinite resolution points of closest approach can be determined with arbitrary precision. However, curves used in object recognition come from observations, where precision is limited by sampling interval, systematic deviations, and random variations. (It should never be assumed that noise can be eliminated by improving the sensor: well established physical principles guarantee that uncertainty will exist at some level of observation, and engineering trade-offs usually result in sensors which return considerable noise.) Curves from observations usually originate from a sensor (or a low level processing routine) as lists of points. Simply choosing closest points on a pair of curves usually selects points which deviate considerably from the 'ideal' curves, and experience has shown that these points are not well suited for matching. Two methods of choosing representative points are investigated here: aggregation of a centroid-like point using a number of points on each curve near the noisy point of closest approach, and determination of the points of closest approach along smooth polynomial fits to both curves in their regions of closest approach.

In the experiments described below these two methods of selecting representative points are investigated for several choices of how much of the curves to use with different amounts of pseudo-normal random noise. The effectiveness of each method was determined by measuring how far the object was from the model after having been translated and rotated to coincide with the model on the basis of parameters returned from the match.

3. Procedures and Programs

The investigations reported here center on the behavior of two methods of selecting points to represent a set of noisy curves, but other procedures have necessarily been developed and used. Since it is difficult to vary the amount of noise in a controlled manner with real measurements, procedures to generate synthetic data and to corrupt this data realistically were developed. The 'matching' procedure [SS85][BSSS86] has been used in a somewhat novel way, and a method of determining the 'separation' of two sets of curves has been developed.

3.1. Data Generation and Corruption

The objects used in these experiments derived from one "view" of a synthetic flower pot with polka-dots. The flower pot was constructed by mapping circles onto the surface of a cone, with distance along the surface of the cone approximated to avoid elliptical integrals. The circles were sampled at regular intervals to provide a list of points which represented the edges of the dots; a "view" of the pot was constructed by keeping only those parts of the circles which lay in the front 80% of the pot after it had been rotated by some angle about its (vertical) axis of symmetry. The pot had 12 circles spaced irregularly about it, and it was possible to reconstruct the pot by repeatedly matching views made at intervals of about 30 degrees using only points of closest approach of the circles as feature points. The individual curves were designed to be nearly indistinguishable: all circles had a radius of 1.6 units except one with a radius of 1.9 units. Each view was stored as a text file consisting of lists of x,y,z coordinates representing the curves, in the same manner as curves observed by the range sensor. When a curve was closed (i.e. a complete circle) this was indicated by setting the last point in the list equal to the first point. Some inaccuracy was introduced in the generation of the data by listing only 3 decimal places, a precision comparable to that obtainable with a very fine depth sensing device.

This nearly perfect data was corrupted by randomization of the starting point of closed figures, randomization of the direction of listing of points, and addition of pseudo-random noise having an something like a normal distribution. For purposes of testing, data sets could be arbitrarily translated and rotated (transformations which the matching program recovered perfectly). The corrupted data was quite similar to measurements typically derived from registered depth and intensity images [BSSS86]. The corrupted data resembled real data also in that the data points did not occur at regular intervals along the curves. To provide a set of points to be used for all processes described below the curves were sampled at intervals of equal arc length. A sampling interval of 0.3 was used in all these experiments, since this was found to be optimum in earlier work [BSSS86] with curves of similar size.

As data base for the random numbers a file of 2550 randomly distributed digits between 0 and 7 was created. The digits 0 through 7 had frequencies of 999, 583, 348, 260, 160, 96, 75, and 29 respectively. Different patterns were created by stepping through the file at different positive and negative intervals (avoiding factors of 2550) and starting at different offsets. Noise was added to each point by adding A*n1 to x, A*n2 to y, and A*n3 to z, where n1-3 are subsequent random digits and A = av.dev/k is a scale factor. The constant k (3.48) was chosen empirically for this set of numbers so the observed RMS deviation coincided in magnitude with the input average deviation. This procedure provided sufficient variety of noise so matching was not spuriously affected, and the corrupted curves visually resembled actual data. The amount of noise used varied from practically nothing to several times that typically encountered with real data.

3.2. Selection of representative points

Two methods were used to determine representative "points of closest approach": in one the points were formed as the centroids of all points within a certain arc-length of the initial point of closest approach, while in the other the closest points along quadratic or cubic fits to all points within a certain distance from the initial point of closest approach were used. Initial points of closest approach were simply the pair of points on the two curves where were closest to each other. Here (as in all further work) the points used to describe the curves were those sampled at equal intervals along the curves.

To form the centroid-like representative point coordinate values for the 2n sampled points adjacent to the initial point (n on each side) were averaged. If the curve was closed all 2n points were used, but if the curve was open and a point of closest approach was less than n from an end fewer points were used, the reduction being equal on both sides and for both object and model curves. The number n was varied in the course of the experiments, taking values 5, 9, and 15, this last value corresponding to most of the data in these figures.

Low order polynomial approximations of coordinate values vs arc length were made in the region of the initial points using a least squares procedure. Constants in the equations

$$X = a + bs + cs^2 \qquad \text{and} \qquad X = a + bs + cs^2 + ds^3$$

were found for X = x, y, or z. Separate constants were determined for each coordinate, thus a curve was described by 9 constants in the quadratic case case and by 12 in the cubic approximation. Here s is arc length from the initial point. All consecutive points less than a particular maximum arc length in each direction were used in the fits; maximum arc length values used in the experiments were 1.5, 2.0, 2.7, and 4.5, corresponding approximately to the values of n used in the centroid method. (A second parametrization was tested in which s was the direct distance from the starting point rather than the distance along the curve, and the experimental results were very similar to those of the arc length parametrization.) Constants were determined for each coordinate and each curve, then the point of closest approach of the two smooth parametrized curves was located numerically (within an epsilon of 0.003 along each curve).

Since performance of these procedures to choose representative points was measured by matching complete sets of curves, the choice of which pairs of curves should be used to represent the set could also be important. In this work the selection of pairs of curves was made by the experimenter with knowledge of which curves on the observed object corresponded to which curves on the model. Matches with many different sets of pairs were made for better indication of the performance of the algorithms under a variety of conditions.

3.3. Matching of sets of representative points

After the observed and model sets of representative points were found the routine used to discover the best match of curves in three-space [BSSS86] was used to return the minimum least squares distance which obtains between the two sets of points, as well as the transformation (translation and rotation) which must be applied to the observed set of points to bring it into best juxtaposition with the set of model points. The algorithm requires a minimum of 8 points, and if fewer curves were used additional points were interpolated between the representative points. There was no requirement that the points be ordered in any way except that model and observed points be listed in the same order.

3.4. Evaluation of matching: separation of two sets of curves

Because a very limited number of points was used, the least squares distance returned by the matching routine does not reflect the degree to which the two sets of curves correspond after the observed curve has been translated and rotated to over-lie the model curve. Fig. 2 suggests the problem: how far is the set of heavy curves from the light curves. The solution is relatively simple when the exact point-to-point correspondence between all curves is known, but this is usually not the case, particularly when curves are corrupted by noise. The central ideas of the method presented here are to discover which curve in one set corresponds to which curve in the other, then to follow corresponding curves point by point and accumulate the distance between the curves.

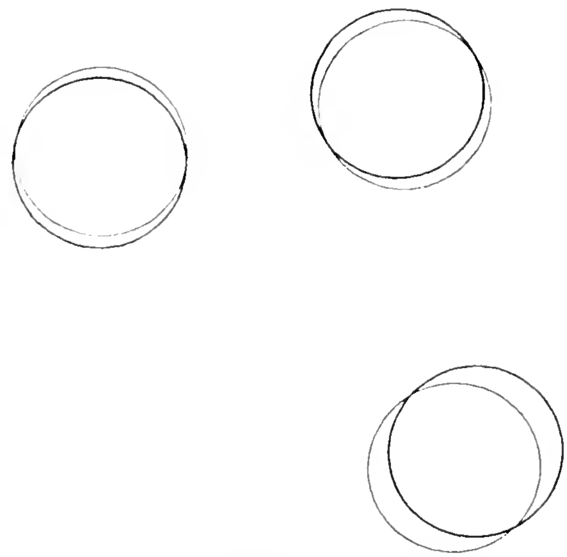

Fig. 2. Heavy curves are slightly rotated with respect to light curves.

Every point of the model (using the sampled points, of course) was placed in a hash table, the hash function being designed to divide space into cubes 0.5 units on a side. Each curve of the observed object was then checked point by point to see if it was close to a point of a model curve using the same hash function, and the identity of the closest pair of points was retained. The hashing procedure was very simple, and placing a point in a cube already occupied overwrote the previous point. Although in principle curves could be very close in space and not intersect in the table (by occupying adjacent cubes) this appeared to cause no difficulty in practice: a multiple hashing function which ensured that any point would find all points within 0.5 units did not improve performance of the procedure for finding corresponding curves.

Using the closest pair of points (one on a model curve, one on an observed curve) as starting points and after discovering the relative listing directions (by checking for closeness of points about five points away from the starting points in both directions), the separation of the two curves was accumulated while following one curve point by point. (If no valid combination of listing directions could be found no separation was accumulated.) The process continued until all observed curves had been checked for hits (i.e. correspondence to a model curve) and followed.

The separation between two corresponding curves A and B was determined by accumulating the squares of the "distances" (determined as below), dividing by the number of points, then taking the root; it is thus an RMS separation. Starting points and listing directions were found as described above. The current point of curve A was incremented and the following two points on curve B were checked to see if either was closer to the current point on A. When a closer point was found the current point B was set to this, and the next following two points were checked. In the absence of noise the two current points (on A and B) would increment nicely in step; with noise the current point B would hang up then jump ahead several points (for obvious reasons). In practice even this two deep look ahead was sometimes insufficient, and an additional constraint that if B failed to increment three times running it was automatically incremented was implemented. The curve following stopped when either of the curves ended, although closed curves were followed in a natural manner. If excessive "distances" were observed for three out of five points, or really excessive "distances" observed once curve following also ended. After a pair of curves was followed to one end the routine returned to the starting points and followed the curves in the opposite directions.

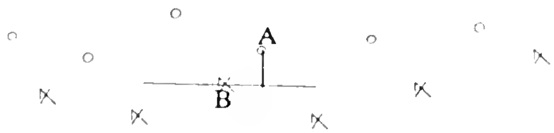

Fig. 3. Heavy line shows the distance between points A and B

The "distance" between the two current points was calculated as the distance between the point on A and the line through the point on B parallel to the line through the two points adjacent to the point on B, as shown in Fig. 3. This procedure was used to reduce the effect of sampling interval, since simply accumulating distance between points could make perfectly juxtaposed smooth curves appear to be separated if they were not sampled at the same points.

Experience and tests suggest that the separation measured using this procedure is meaningful for values up to about .5 but with limited power of discrimination above this, the exact value depending on thresholds in the routine. The reason for this is that curves with greater separations are not followed very far into the regions of greater distances. This is desirable, since otherwise there is little protection against following an incorrect curve. For separations less than 0.5 the value returned appears to give an excellent measure of the separation of two sets curves, although it cannot be used numerically as an absolute determination of the separation of two curves. Separation thus defined was used to measure the quality of the matches described in the experiments below.

4. Experiments and Results

Two algorithms were implemented to choose representative points for pairs of curves, and these were tested using computer generated data sets corrupted by varying amounts of noise. The sets of curves used in all of these experiments consisted of five closed circles on the surface of a cone (analogous to the edges of circular spots on a flower pot) plus somewhat less than half of a large circle about the axis of the cone (analogous to a rim of a flower pot). This rim was used in very few matches reported here, and had negligible effect. Each algorithm was tested by matching three objects (i.e. sets of curves) against each other for several conditions of length of curve and several selections of curve pairs. The quality of match was evaluated on the basis of the separation of the objects after they had been placed in best juxtaposition (using rotation and translation parameters returned by the match).

All three objects were the same basic data set with data translated and randomized. Object 1 was generated by translation 5 units along each of x,y, and z axes followed by randomization of starting point and listing order for each curve; object 2 was not translated but was randomized in a similar but not identical manner; object 3 was translated -5 units along x,y, and z axes and randomized. Noise was added to each object in the same pattern, but since the starting point of the curves had been changed before noise was added similar noise patterns were unlikely to develop in corresponding places on the curves being matched. When noise was reduced to zero matching these objects resulted in perfect juxtaposition and separation of 0.0.

Measuring the separation between two sets of curves was accomplished using the curve following procedure described in section 3.4. This procedure was tested by measuring the separation between two sets of curves which were slightly displaced by several increasing amounts. "Displacements" such as translations, rotations, expansions, skews, and combinations of these all showed increasing separation with increasing amount of displacement for all conditions of noise. Table 1 summarizes separations recorded for 5 different translations with 5 different amounts of noise, and it is typical of all tests of the procedure for measuring separation. The displacements in Table 1 are along x,y, and z axes; the separations are the average of 5 different conditions of noise (having the same average deviation, merely different random pattern).

Transl	Separations for noise of				
	0.01	0.05	0.1	0.15	0.2
0.00	0.068	0.114	0.174	0.275	0.404
0.05	0.132	0.161	0.215	0.297	0.401
0.10	0.228	0.249	0.285	0.362	0.449
0.20	0.317	0.344	0.367	0.434	0.513
0.30	0.350	0.404	0.452	0.498	0.568
0.40	0.425	0.475	0.497	0.567	0.675

Table 1 Test of separation routine for objects translated and with noise added.

The observed separations clearly show increases with both amount of translation and amount of noise. The separation of the combined effects is not the sum of separations due to noise and translation individually because the two effects interfere with each other slightly: at small translations the observed separation is due mostly to noise, while at large translations the noise contribution is relatively unimportant. The tests establish that this procedure returns a useful measure of the separation between two sets of curves, particularly when they have similar amounts of noise.

It should be noted that this test situation is nearly optimum in that the routine was able to follow curves their entire length. When excessively large displacement exists for some portion of two curves it is not recorded, resulting in artificially low separation readings. In practice observed separations tend to peak out at something over 0.5, the exact value depending on threshold values within the routine which define excessive distance between points.

The algorithm for measuring separation of two curves does not act symmetrically, that is, it can make a difference which of the two curves is is followed point by point. The difference is largest when noise is added to only one curve, since in one case points closer to the followed curve will be consistently chosen. In the experiments reported in Table 1 noise was added to only one curve, and separations were consistently different from the average of both directions (the value reported in the table) by ±20%. All separations reported below are averages of separations measured in both directions, but the two values rarely differed by more than 5% when both sets of curves were similarly noisy. The matching routine itself acts perfectly symmetrically, and it makes no difference which object is considered to be the model and which the observed.

Quadratic and cubic polynomials expressing coordinates vs arc length were generated using a least squares procedure for four different lengths of the noisy curves. After the coefficients for the smooth curves were determined the closest points on these fits were found (within an epsilon of 0.003 units along the curves) and used as representative points. The different conditions are identified in Table 2 using Qu for quadratic and Cu for cubic, and the arc length on both sides of the starting point which was used: 1.5, 2.0, 2.7, 4.5.

The centroid procedure used three different amounts of the curves, measured in the number of points from the starting point on both sides. Since the curves were sampled at intervals of 0.3, the numbers 5, 9, and 15 correspond to arc lengths of 1.5, 2.7, and 4.5 (on each side), comparable to the lengths used in the fitting procedures. In Table 2 the conditions are identified using Ce (for centroid) and the number of points used in the aggregation of the centroid.

Table 2 presents the quality of match observed on the average for the various conditions tested with the two procedures. The objects matched were identical except for translations, randomization of starting points and listing order within the component curves, and addition of random noise in varying amounts. High quality of match corresponds to small separation of the curves after they have been brought into closest juxtaposition using the match of their representative points. Separations for the polynomial cases are averages of 5 different sets of two pairs of points, while separations for centroid cases are averages of 9 different sets of two pairs of points. This difference is not significant. Averages for matches of three different pairs of objects are included to indicate typical variation.

Type	Objects	Separations for noise of average deviation					
		0.001	0.005	0.01	0.02	0.05	0.1
Qu 1.5	1 2	0.213	0.224	0.212	0.223	0.212	0.262
	1 3	0.206	0.214	0.216	0.214	0.218	0.260
	2 3	0.261	0.248	0.256	0.241	0.221	0.293
Qu 2.0	1 2	0.306	0.286	0.265	0.279	0.301	0.350
	1 3	0.311	0.291	0.272	0.268	0.286	0.246
	2 3	0.334	0.335	0.305	0.331	0.289	0.364
Qu 2.7	1 2	0.341	0.348	0.323	0.366	0.416	0.417
	1 3	0.252	0.264	0.254	0.262	0.294	0.346
	2 3	0.367	0.372	0.384	0.385	0.366	0.374
Cu 1.5	1 2	0.063	0.053	0.072	0.081	0.156	0.238
	1 3	0.070	0.056	0.078	0.079	0.146	0.253
	2 3	0.069	0.070	0.070	0.094	0.131	0.239
Cu 2.0	1 2	0.113	0.094	0.130	0.115	0.174	0.245
	1 3	0.112	0.081	0.125	0.119	0.169	0.208
	2 3	0.102	0.105	0.114	0.122	0.128	0.273
Cu 2.7	1 2	0.194	0.137	0.127	0.129	0.142	0.229
	1 3	0.225	0.154	0.154	0.114	0.170	0.184
	2 3	0.202	0.195	0.183	0.185	0.163	0.245
Cu 4.5	1 2	0.453	0.539	0.543	0.565	0.396	0.340
	1 3	0.544	0.601	0.595	0.541	0.327	0.295
	2 3	0.677	0.595	0.670	0.585	0.245	0.275
Ce 5	1 2	0.097	0.105	0.117	0.244	0.237	0.263
	1 3	0.090	0.095	0.108	0.105	0.210	0.233
	2 3	0.101	0.068	0.094	0.097	0.189	0.286
Ce 9	1 2	0.066	0.064	0.070	0.146	0.181	0.229
	1 3	0.061	0.066	0.074	0.074	0.160	0.211
	2 3	0.076	0.049	0.070	0.077	0.163	0.259
Ce 15	1 2	0.027	0.030	0.036	0.050	0.102	0.178
	1 3	0.028	0.028	0.031	0.039	0.097	0.183
	2 3	0.037	0.024	0.044	0.063	0.115	0.229

Table 2 Quality of Matches

The matches using quadratic approximations (coordinates vs arc length) are very bad, and in fact not useful, while using the cubic approximation to smooth noisy curves results in good matches when only a small section of the curve is used. As the amount of curve used in the fit is increased the quality of fit decreases, apparently because the location of points of closest approach of smooth curves is quite sensitive to slight changes of shape. The cubic fit procedure is superior to the centroid procedure when a small amount of the curves are used because, while the starting points for both procedures are the same, the fitting algorithm allows the representative point to move from the starting point more easily. (The starting point is almost always a point of considerable deviation from the noise-free curve.)

However, the centroid-like representative points provide continually better matches as the number of data points used increases. Increasing noise tends to decrease the quality of match under all circumstances; the quadratic fit procedure appears to be relatively less affected by noise because it behaves so poorly when there is little noise. It should be observed that since the noise-free curves have a length of about 10 in these units, the polynomial fits Qu 4.5 and Cu 4.5 used almost all of a curve for the fit.

The above matching results were obtained using two pairs of representative points for each object. Since the objects consist of six curves, there could be as many as 15 pairs of points; the five closed curves guarantee at least 10 pairs. Table 3 shows the effect that using different numbers of pairs of representative points has on the quality of matches using the centroid algorithm with different amounts of data (as in Table 2). The separations given are average values: for 2 pairs, 9 different combinations are averaged; for 3 pairs, 5; for 4, 4; for 5, 2; for 10 pairs, only one combination was used. In Table 3 separations are given for matching objects 1 and 3 only, but separations for the other matches (1 with 2, 2 with 3) are similar in size and behavior.

No.Pts	No.Prs	Separations for average deviations of					
		0.001	0.005	0.01	0.02	0.05	0.1
5	2	0.090	0.095	0.108	0.105	0.210	0.233
5	3	0.067	0.068	0.085	0.090	0.136	0.222
5	4	0.052	0.064	0.085	0.141	0.149	0.228
5	5	0.040	0.047	0.067	0.081	0.163	0.234
5	10	0.031	0.048	0.043	0.082	0.101	0.165
9	2	0.061	0.066	0.074	0.074	0.160	0.211
9	3	0.049	0.045	0.068	0.068	0.117	0.208
9	4	0.039	0.046	0.059	0.095	0.122	0.198
9	5	0.033	0.038	0.049	0.064	0.132	0.209
9	10	0.023	0.038	0.033	0.057	0.092	0.166
15	2	0.028	0.028	0.031	0.039	0.097	0.183
15	3	0.026	0.027	0.035	0.041	0.096	0.188
15	4	0.021	0.026	0.028	0.042	0.095	0.172
15	5	0.020	0.022	0.027	0.041	0.098	0.182
15	10	0.021	0.026	0.027	0.035	0.088	0.161

Table 3 Centroid matches for objects 1 & 3

Increasing the number of representative points used in a match improves the result (as shown by the smaller separations recorded), and also (as noted above) increasing the number of points used to generate each representative point improves a match. It is interesting to compare the separations on lines 15-3 and 9-5 (for example): both have used a total of about 180 sampled points of the data, yet the separations of the 15-3 line are significantly smaller. It thus appears more important to use high quality representative points than to use a larger number of them.

5. Summary

The most significant result of these experiments is that sets of curves can be matched well using representative points. This should lead to more efficient matching procedures, since curves consisting of several hundred points can be identified using ten representative points. While representative points do not fully describe a set of curves they provide points useful for matching with similarly processed model curves.

The use of centroid-like points to represent curves appears to be superior to using polynomial fits because the quality of the representative point can be more easily improved by using more data points. The weakness of the centroid method is that the starting point (the data point closest to the other curve of the pair) can be easily displaced from the point of closest approach on the "ideal curve" by noise, and considerable improvement in the generation of representative points might be made by attacking this problem. The aggregation of the centroid of a small (or even large) number of points is computationally cheap.

While the cubic fit performed better than the centroid method with a small number of data points, the extension of these results (made with curves which were nearly circles) to more general curves is questionable. It seems likely that higher order polynomials would be needed (with corresponding increase in computing time), particularly in view of the failure of representative points made using quadratic fits.

The successful use of representative points to match sets curves has some implications for further work in object recognition. (Recall that procedures by which curves significant to the identity of an object can be extracted from registered range and intensity images are being developed.) Clearly the representative points used in these experiments cannot be defined for all pairs of curves, and additional classes of points should be developed. Use of such points for matching objects requires that the points of object and model be listed in corresponding order, and procedures to accomplish this should be developed. (If these procedures are strong enough the actual matching is needed only in rare cases, and perhaps to return the translation and rotation needed to bring the observed object into standard orientation.) It is likely that relationships between representative points will be useful in these procedures.

Many procedures developed in the course of this work are useful utilities, but are of little scientific importance. The routines used to generate synthetic data, the general purpose filter for the modification of data, and the generation of pseudo-normal random noise all fall into this category. However, the procedure for determining which curves are close in space has some interesting possible applications. This use of hashing to efficiently detect closeness is not limited to curves; the same technique can be used for surfaces or even solids, and in any number of dimensions, since all that is needed is a set of points to represent the curve, surface, etc.

One possible application of this technique is real-time collision avoidance for robots, especially in situations where the motion of a robot is determined by sensory input from a changing environment. In simplest form, the work space would be described by a set of points chosen at some sufficiently fine spacing whose coordinates would be hashed into table. A similar set of points connected to the robot would be checked for collision (using the same hash function) at some time far enough ahead of robot motion that stopping would be possible. Refinements could include variation of the surface sampling mesh and hash cube size depending on speed and distance from obstacles and real-time updating of obstacles, including using multiple robots in the same work-space. It should be noted that computations of actual geometric distance need not be done, as observation of closeness in any coordinate would provide sufficient warning.

References

[B84] M. Brady, "Representing Shape," in *Proc. Int. Conf. on Robotics* (Atlanta, 1984, IEEE), pp. 256-265.

[BSSS86] C.M Bastuscheck, E. Schonberg, J.T. Schwartz and M. Sharir, "Object recognition by three-dimensional curve matching", *Int. J. of Intelligent Systems*, **1**, pp. 105-132, 1986.

[BC82] R.C. Bolles and R.A. Cain, "Recognizing and locating partially visible objects: the local-feature-focus method," *Inter. J. Robotics Research*, **1-3** pp. 57-82, 1982.

[BJ85] P.J. Besl and R.C. Jain, "Three-dimensional object recognition," *ACM Computing Surveys* **17**, 75-145 (1985).

[H85] M. Herman, "Generating detailed scene descriptions from range images," in *Proc. 1985 IEEE Int. Conf. on Robotics and Automation*, 1985, pp. 426-431.

[FH83] O. D. Faugeras and M. Hebert, "A 3-D recognition and positioning algorithm using geopmetrical matching between surface primitives," in *Proc. 8th IJCAI*, 1983, pp. 996-1002.

[HB84] P. Horaud and R. Bolles, "3DPO's strategy for matching three-dimensional objects in range data," in *Proc. Inter. Conf. on Robotics*, IEEE, New York, 1984, pp. 78-85.

[OS81] M. Oshima and Y. Shirai, "Object recognition using three-dimensional information," *IEEE Trans. on Pattern Analysis and Machine Intelligence* **PAMI-5**, 353-361 (1983) also in *Proc. 7th IJCAI*, 1981, pp. 601-606.

[SS71] Y. Shirai and M. Suwa, "Recognition of polyhedrons with a range finder," in *Proc. 2nd IJCAI*, 1971, pp. 80-87.

DATE DUE
